

9/981,613

[> home](#) [> about](#) [> feedback](#) [> login](#)

US Patent &amp; Trademark Office

Try the *new* Portal design

Give us your opinion after using it.

## Search Results

Search Results for: [partition\* and (table? or database?) and ((row? or tuple? or record?)  
<paragraph> (id? or key? or (primary <near> key?))) and value? and "64" and \*bit? ]  
Found 73 of 127,944 searched.

## Search within Results

[> Advanced Search](#)[> Search Help/Tips](#)Sort by: [Title](#) [Publication](#) [Publication Date](#) [Score](#) [Binder](#)Results 1 - 20 of 73 [short listing](#)

1

2

3

4

1 [High-performance sorting on networks of workstations](#)

99%



Andrea C. Arpaci-Dusseau , Remzi H. Arpaci-Dusseau , David E. Culler , Joseph M. Hellerstein , David A. Patterson

**ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data June 1997**

Volume 26 Issue 2

We report the performance of NOW-Sort, a collection of sorting implementations on a Network of Workstations (NOW). We find that parallel sorting on a NOW is competitive to sorting on the large-scale SMPs that have traditionally held the performance records. On a 64-node cluster, we sort 6.0 GB in just under one minute, while a 32-node cluster finishes the Datamation benchmark in 2.41 seconds. Our implementations can be applied to a variety of disk, memory, and processor configura ...

2 [Adaptive algorithms for set containment joins](#)

97%













Sergey Melnik , Hector Garcia-Molina

**ACM Transactions on Database Systems (TODS) March 2003**

Volume 28 Issue 1

A set containment join is a join between set-valued attributes of two relations, whose join condition is specified using the subset (&sube;) operator. Set containment joins are deployed in many database applications, even those that do not support set-valued attributes. In this article, we propose two novel partitioning algorithms, called the Adaptive Pick-and-Sweep Join (APSJ) and the Adaptive Divide-and-Conquer Join (ADCJ), which allow computing set containment joins efficiently. We show that ...


- 3 IS '97: model curriculum and guidelines for undergraduate degree programs in information systems 93%  
 Gordon B. Davis , John T. Gorgone , J. Daniel Couger , David L. Feinstein , Herbert E. Longenecker  
**ACM SIGMIS Database , Guidelines for undergraduate degree programs on Model curriculum and guidelines for undergraduate degree programs in information systems**  
December 1997  
Volume 28 Issue 1
- 4 Affinity-based management of main memory database clusters 88%  
 Minwen Ji  
**ACM Transactions on Internet Technology (TOIT)** November 2002  
Volume 2 Issue 4  
We study management strategies for main memory database clusters that are interposed between Internet applications and back-end databases as content caches. The task of management is to allocate data across individual cache databases and to route queries to the appropriate databases for execution. The goal is to maximize effective cache capacity and to minimize synchronization cost. We propose an affinity-based management system for main memory database cLUsters (*ALBUM*). *ALBUM* executes ea ...
- 5 IP lookups using multiway and multicolumn search 88%  
 Butler Lampson , Venkatachary Srinivasan , George Varghese  
**IEEE/ACM Transactions on Networking (TON)** June 1999  
Volume 7 Issue 3
- 6 Logical and physical design issues for smart card databases 87%  
 Cristiana Bolchini , Fabio Salice , Fabio A. Schreiber , Letizia Tanca  
**ACM Transactions on Information Systems (TOIS)** July 2003  
Volume 21 Issue 3  
The design of very small databases for smart cards and for portable embedded systems is deeply constrained by the peculiar features of the physical medium. We propose a joint approach to the logical and physical database design phases and evaluate several data structures with respect to the performance, power consumption, and endurance parameters of read/program operations on the Flash-EEPROM storage medium.
- 7 Rendering and simulation: Low latency photon mapping using block hashing 87%  
 Vincent C. H. Ma , Michael D. McCool  
**Proceedings of the ACM SIGGRAPH/EUROGRAPHICS conference on Graphics hardware** September 2002  
For hardware accelerated rendering, photon mapping is especially useful for simulating caustic lighting effects on non-Lambertian surfaces. However, an efficient hardware algorithm for the computation of the k nearest neighbours to a sample point is required. Existing algorithms are often based on recursive spatial subdivision techniques, such as kd-trees. However, hardware implementation of a tree-based algorithm would have a high latency, or would require a large cache to avoid this latency on a ...
- 8 Fast algorithms for universal quantification in large databases 87%

-  Goetz Graefe , Richard L. Cole  
**ACM Transactions on Database Systems (TODS)** June 1995  
Volume 20 Issue 2  
Universal quantification is not supported directly in most database systems despite the fact that it adds significant power to a system's query processing and inference capabilities, in particular for the analysis of many-to-many relationships and of set-valued attributes. One of the main reasons for this omission has been that universal quantification algorithms and their performance have not been explored for large databases. In this article, we describe and compare three known algorithms ...
- 9 Join processing in relational databases 87%  
 Priti Mishra , Margaret H. Eich  
**ACM Computing Surveys (CSUR)** March 1992  
Volume 24 Issue 1  
The join operation is one of the fundamental relational database query operations. It facilitates the retrieval of information from two different relations based on a Cartesian product of the two relations. The join is one of the most difficult operations to implement efficiently, as no predefined links between relations are required to exist (as they are with network and hierarchical systems). The join is the only relational algebra operation that allows the combining of related tuples from ...
- 10 WCRC: An ANSI SPARC machine architecture for data base management 85%  
 Sudhir K. Arora , S. R. Dumpala , K. C. Smith  
**Proceedings of the 8th annual symposium on Computer Architecture** May 1981  
Several data base machine architectures have been proposed in the past few years. The next generation of these machines must simultaneously support different data models on the same physical data. One approach to this problem, GDBMS, has recently been described in the literature. This paper presents the architecture of another approach: the well-connected relation computer (WCRC). The architectures of these two computers are compared, and a preliminary performance evaluation showing that WC ...
- 11 Data page layouts for relational databases on deep memory hierarchies 85%  
 Anastassia Ailamaki , David J. DeWitt , Mark D. Hill  
**The VLDB Journal &mdash; The International Journal on Very Large Data Bases**  
November 2002  
Volume 11 Issue 3  
Relational database systems have traditionally optimized for I/O performance and organized records sequentially on disk pages using the N-ary Storage Model (NSM) (a.k.a., slotted pages). Recent research, however, indicates that cache utilization and performance is becoming increasingly important on modern platforms. In this paper, we first demonstrate that in-page data placement is the key to high cache performance and that NSM exhibits low cache utilization on modern platforms. Next, we ...
- 12 A simple bounded disorder file organization with good performance 85%  
 David B. Lomet  
**ACM Transactions on Database Systems (TODS)** October 1988  
Volume 13 Issue 4  
A bounded-disorder (BD) file is one in which data are organized into nodes that are indexed,

e.g., by means of a B-tree. The data nodes are multibucket nodes that are accessed by hashing. In this paper we present two important improvements to the BD organization as originally described. First, records in a data node that overflow their designated primary bucket are stored in a single overflow bucket which is itself a bucket of the data node. Second, when file space needs to be increased, pa ...

13 Research track: Information-theoretic co-clustering

84%


 Inderjit S. Dhillon , Subramanyam Mallela , Dharmendra S. Modha

**Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining** August 2003

Two-dimensional contingency or co-occurrence tables arise frequently in important applications such as text, web-log and market-basket data analysis. A basic problem in contingency table analysis is *co-clustering: simultaneous clustering* of the rows and columns. A novel theoretical formulation views the contingency table as an empirical joint probability distribution of two discrete random variables and poses the co-clustering problem as an optimization problem in *information theory*

14 B+ retake: sustaining high volume inserts into large data pages

84%


 Kurt W. Deschler , Elke A. Rundensteiner

**Proceedings of the 4th ACM international workshop on Data warehousing and OLAP** November 2001

Modern ad-hoc data mining queries often run on databases over a terabyte in size. At this scale, large data pages are required to obtain sufficient disk performance. Unfortunately, these large data pages greatly increase update costs, especially for packed structures such as the B+ tree. In a frequently updated warehouse, users are often forced to decide between query performance and update performance in order to meet maintenance time windows. Solutions that provide both are welcome. In this pap ...

15 Partial-match retrieval using hashing and descriptors

84%


 K. Ramamohanarao , James A. Thom , John W. Lloyd

**ACM Transactions on Database Systems (TODS)** December 1983  
Volume 8 Issue 4

This paper studies a partial-match retrieval scheme based on hash functions and descriptors. The emphasis is placed on showing how the use of a descriptor file can improve the performance of the scheme. Records in the file are given addresses according to hash functions for each field in the record. Furthermore, each page of the file has associated with it a descriptor, which is a fixed-length bit string, determined by the records actually present in the page. Before a page is accessed to s ...

16 Fast string searching in secondary storage: theoretical developments and experimental results


84%

 Paolo Ferragina , Roberto Grossi

**Proceedings of the seventh annual ACM-SIAM symposium on Discrete algorithms**  
January 1996

17 Balanced multidimensional extendible hash tree

82%

 Ekow J. Otoo

**Proceedings of the fifth ACM SIGACT-SIGMOD symposium on Principles of database**

## systems June 1985

- 18 Peer to peer networks: Tarzan: a peer-to-peer anonymizing network layer 82%  
Michael J. Freedman , Robert Morris  
**Proceedings of the 9th ACM conference on Computer and communications security**  
November 2002  
Tarzan is a peer-to-peer anonymous IP network overlay. Because it provides IP service, Tarzan is general-purpose and transparent to applications. Organized as a decentralized peer-to-peer overlay, Tarzan is fault-tolerant, highly scalable, and easy to manage. Tarzan achieves its anonymity with layered encryption and multi-hop routing, much like a Chaumian mix. A message initiator chooses a path of peers pseudo-randomly through a restricted topology in a way that adversaries cannot easily influence ...
- 19 Optimal VLSI circuits for sorting 82%  
Richard Cole , Alan Siegel  
**Journal of the ACM (JACM)** October 1988  
Volume 35 Issue 4  
This work describes a large number of constructions for sorting  $N$  integers in the range  $[0, M - 1]$ , for  $N \leq M \leq N^2$ , for the standard VLSI bit model. Among other results, we attain: VLSI sorter constructions that are within a constant factor of optimal size, for all  $M$  and almost all running times  $T$ . a ...
- 20 Office Information Systems and Computer Science 82%  
Clarence A. Ellis , Gary J. Nutt  
**ACM Computing Surveys (CSUR)** January 1980  
Volume 12 Issue 1

---

Results 1 - 20 of 73    short listing

◀  
Prev  
Page

1   2   3   4

Next  
Page ▶

---

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.